

UNITED STATES PATENT APPLICATION FOR
DETERMINING ACCURACY OF A CLASSIFIER

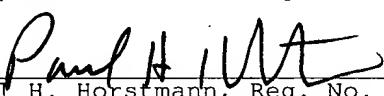
Inventors:

Henri Jacques Suermondt
Tom Elliott Fawcett

CERTIFICATE OF MAILING BY "EXPRESS MAIL"
UNDER 37 C.F.R. § 1.10

"Express Mail" mailing label number: EJ138444457US
Date of Mailing: 9-17-01

I hereby certify that this correspondence is
being deposited with the United States Postal
Service, utilizing the "Express Mail Post Office to
Addressee" service addressed to **Assistant
Commissioner for Patents, Washington, D.C. 20231** and
mailed on the above Date of Mailing with the above
"Express Mail" mailing label number.


Paul H. Horstmann, Reg. No. 36,167
Signature Date: 9-17-01

BACKGROUND OF THE INVENTION

Field of Invention

5 The present invention pertains to the field of classifiers. More particularly, this invention relates to determining accuracy of a classifier.

Art Background

10 A classifier may be defined as an entity that associates an item with one or more of a set of categories. Typically, a classifier determines which categories with which to associate an item in response to attributes of the item. Examples of classifiers are numerous and include rule-based
15 systems, neural networks, Bayesian probability systems, as well as human beings.

20 It is commonly desirable to determine the accuracy of a classifier. For example, the accuracy of a classifier may be used to improve the classifier or to compare the relative accuracy of different classifiers.

25 One prior method for determining the accuracy of a classifier is to compare classifications rendered by the classifier with classifications rendered by an expert. Such a method usually yields a boolean correct/incorrect indication of accuracy with respect to the classification of individual items. The
30 boolean indications may be aggregated to obtain a measure of the accuracy of the classifier. Unfortunately, such correct/incorrect indications are usually of relatively limited utility in evaluating

the relative performance of a classifier. Moreover, such indications usually do not provide information on the extent to which an individual classification is wrong.

10007909

SUMMARY OF THE INVENTION

A method for determining accuracy of a classifier is disclosed which provides an indication of the degree of correctness of the classifier rather than a mere correct/incorrect indication. The accuracy of the classifier is determined by determining a set of categories of an arrangement of categories selected for an item by the classifier and 10 determining a set of categories of the arrangement selected for the item by an authoritative classifier. An accuracy measure which indicates a degree of correctness of the classifier is then determined based on the categories selected by the classifier 15 and the categories selected by the authoritative classifier.

Other features and advantages of the present invention will be apparent from the detailed 20 description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is described with respect to particular exemplary embodiments thereof and reference is accordingly made to the drawings in which:

5 **Figure 1** shows a method for determining an accuracy of a classifier according to the present
10 teachings;

15 **Figure 2** shows an accuracy evaluator that determines an accuracy of a classifier according to the present teachings;

20 **Figure 3** shows a method for determining an accuracy measure based on the categories selected by a classifier and the categories selected by an authoritative classifier;

25 **Figure 4** shows an example arrangement of a set of categories A-E into which an item may be placed;

30 **Figure 5** shows the result of step 110 for the example categories A-E and an example classification rendered by an authoritative classifier;

35 **Figure 6** shows the result of step 112 for the example categories A-E and an example classification rendered by a classifier;

Figure 7a-7c show an example arrangement of a set of categories A-K into which an item may be placed;

- 5 **Figure 8** illustrates the training of a classifier in response to an accuracy determined according to the present teachings.

CONFIDENTIAL - SECURITY INFORMATION

DETAILED DESCRIPTION

5 **Figure 1** shows a method for determining an accuracy of a classifier according to the present teachings. At step 100, a set of zero or more categories selected for an item by the classifier is determined. The categories are selected by the classifier from among an arrangement of possible categories into which the item may be placed. The 10 arrangement of categories may be, for example, a hierarchy of categories.

15 At step 102, a set of one or more categories selected for the item by an authoritative classifier is determined. The categories are selected by the authoritative classifier from among the arrangement of possible categories. At step 104, an accuracy measure is determined based on the categories selected by the classifier and the categories selected by the authoritative classifier. The 20 accuracy measure provides an indication of the accuracy of the classifier for the item.

25 **Figure 2** shows an accuracy evaluator 16 that determines an accuracy 18 of a classifier 12 according to the present teachings. The accuracy 18 indicates the relative distance between the categories selected by the classifier 12 for an item 10 and the categories selected for the item 10 by an 30 authoritative classifier 14.

The item 10 may be anything that may be classified. Examples include products, human beings,

documents, natural elements, etc. The item 10 may be classified by the classifier 12 and the authoritative classifier 14 using one or more attributes of the item 10.

5

The classifier 12 may be any type of classifier examples of which include rule based systems including automated systems, neural networks, human experts, etc. The classifier 12 may classify items 10 into zero or more classes or categories which may be organized in a hierarchy such that if an item is in a class it is also considered in the parent of that class.

15

The authoritative classifier 14 may be a human expert or a highly accurate automated classifier.

20

The accuracy evaluator 16 may be a human being or group of human beings or an automated system that performs the present methods.

25

Figure 3 shows a method for determining an accuracy measure based on the categories selected by the classifier 12 for the item 10 and the categories selected for the item 10 by an authoritative classifier 14. The method steps shown are performed by the accuracy evaluator 16 in response to the categories selected by the classifier 12 and the categories selected by the authoritative classifier 30 14.

At step 110, the accuracy evaluator 16 assigns a true indication to each category in the arrangement

selected by the authoritative classifier 14 for the item 10 and assigns a false indication to each category in the arrangement not selected by the authoritative classifier 14 for the item 10. If a
5 given category is selected by the authoritative classifier 14 for the item 10 then the given category and all of its ancestors are assigned the true indication at step 110.

10 At step 112, the accuracy evaluator 16 assigns a positive indication to each category in the arrangement selected by the classifier 12 for the item 10 and a negative indication to each category in the arrangement not selected by the classifier 12 for
15 the item 10. If a given category is selected by the classifier 12 for the item 10 then the given category and all of its ancestors are assigned the positive indication at step 112.

20 At step 114, the accuracy 18 of the classifier 12 is determined by combining the true, false, positive, and negative indications. The accuracy 18 may have a specified range such as between 0 and 1 or between 0 and 100, etc. For example, a higher value
25 for the accuracy 18 indicates a relatively higher efficacy of the classifier 12 in classifying the item 10.

30 The accuracy 18 in one embodiment is based on a set of measures derived from the true, false, positive, and negative indications. The measures include a true positive count (TP), a false positive count (FP), and a false negative count (FN). The

true positive count may be determined by counting the categories which are assigned the true and positive indications. The false positive count may be determined by counting the categories which are assigned the false and positive indications. The 5 false negative count may be determined by counting the categories which are assigned the true and negative indications.

10 The measures derived from the true, false, positive, and negative indications may include a true negative count (TN). The true negative count may be determined by counting the categories which are assigned the false and negative indications.

15 One or more of the categories into which the item 10 may be classified may be assigned an indifferent indication (I) by the authoritative classifier 14. The categories with the indifferent indication do not contribute to the measures of accuracy. For example, the categories having an indifferent indication are not counted when determining the TP, FP, FN, or TN counts.

20 The accuracy 18 may be obtained by combining an over-conservativeness measure (OC) and an over-aggressiveness measure (OA). The over-conservativeness is the tendency of the classifier 12 to not put an item in enough classes or not put the item deep enough in the hierarchy given the attributes of the item. The over-aggressiveness is the tendency of the classifier 12 to put an item in

more classes or in classes deeper in the hierarchy than is warranted by the attributes of the item.

In one embodiment, the over-conservativeness measure equals the $FN/(TP+FN)$ and the over-aggressiveness measure equals $FP/(FP+TP)$. The over-conservativeness and over-aggressiveness measures may be averaged as follows to obtain the accuracy 18.

$$\text{accuracy} = 1 - \frac{OA + OC}{2}$$

The over-conservativeness and over-aggressiveness measures may be combined using a harmonic mean as follows.

$$\text{accuracy} = 1 - \frac{2}{\frac{1}{OA} + \frac{1}{OC}}$$

where $\text{accuracy}=1$ if $OA=0$ or $OC=0$.

Alternatively, the over-conservativeness and over-aggressiveness measures may be combined as follows.

$$\text{accuracy} = 1 - (\alpha * OA + \beta * OC)$$

The coefficients α and β may be adjusted to adjust the relative importance of the over-aggressiveness and the over-conservativeness in determining the accuracy 18 of the classifier 12. In one embodiment, $\alpha + \beta = 1$.

The TP, FP, FN, TN, OA, OC, and/or combined accuracy measures may be aggregated over multiple test items to obtain a comprehensive measure of the accuracy of the classifier 12. The test items may be
5 any subset of a predetermined test item set.

The accuracy 18 provides an indication of a degree of correctness in the classification rendered by the classifier 12 where the correctness is
10 provided by the authoritative classifier 14. The accuracy 18 may be viewed as providing a measure of distance between the zero or more categories selected by the classifier 12 for the item 10 and the one or more categories selected by the authoritative
15 classifier 14 for the item 10.

Accuracy values determined for different classifiers using the present teachings may be used to evaluate the relative efficacy of the classifiers.
20 Similarly, accuracy values determined before and after changes in the classifier 12 may be used to evaluate the relative goodness or badness in the changes made to the classifier 12.

25 **Figure 4** shows an example arrangement of a set of categories A-E into which the item 10 may be placed. In this example, the categories A-E are arranged in a hierarchy. The categories B and C are children of the category A and the categories D and E are children of the category B. In other examples, the possible categories into which the item 10 may be placed may be discrete categories.
30

In the following example, the authoritative classifier 14 selects the category E for the item 10 whereas the classifier 12 selects the category D for the item 10.

5

Figure 5 shows the result of step 110 for the example categories A-E and the example classification rendered by the authoritative classifier 14 which specifies category E. The category E is marked with the true indication (T) at step 110. In addition, all of the ancestors of category E, the categories A and B in this example, are marked with the true indication at step 110. The remaining categories C and D are marked with the false indication (F) at step 110.

In general, if a given category is an ancestor of multiple categories that are marked with the true indication then the given category is only marked once with the true indication.

Figure 6 shows the result of step 112 for the example categories A-E and the example classification rendered by the classifier 12 which specifies the category D. The category D is marked with the positive indication (P) at step 112. In addition, all of the ancestors of category D, the categories B and A in this example, are marked with the positive indication at step 110. The remaining categories C and E are marked with the negative indication (N) at step 112.

In general, if a given category is an ancestor of multiple categories that are marked with the positive indication then the given category is only marked once with the positive indication.

5

In the example shown in **Figure 6**, TP=2 (categories A and B), FP=1 (category D), TN=1 (category C), FN=1 (category E), OC=1/3, and OA=1/3.

10

15

20

Figure 7a shows an example arrangement of a set of categories A-K into which an item 10 may be placed. In this example, the categories A-K are arranged into a forest, i.e. a collection of trees (hierarchies). Two of these trees consist of singleton nodes (node F and node K). A singleton node is a special case of a very simple tree. The categories B and C are children of category A and the categories D and E are children of the category B. The categories H and I are children of the category G, and the category J is a child of the category I.

25

In the following example, the authoritative classifier 14 selects the categories B, H, and K for the item 10 whereas the classifier 12 selects the categories A, F, H, and K for the item 10. Moreover, the authoritative classifier 14 specifies that it is indifferent about category F.

30

Figure 7b shows the result of step 110 for the example categories A-K and the example classification rendered by the authoritative classifier 14 which specifies categories B, H, and K, and specifies indifferent for category F. The categories B, H, and

K are marked with the true indication (T) at step 110. In addition, all of the ancestors of categories B, H, or K, the categories A and G in this example, are marked with the true indication at step 110. In 5 addition, category F is marked with the indifferent indication (I) in a sub-step of step 110.

Figure 7c shows the result of step 112 for the example categories A-K and the example classification 10 rendered by the classifier 12 which specifies the categories A, F, H, and K. The categories A, F, H, and K are marked with the positive indication (P) at step 112. In addition, all the ancestors of categories A, F, H, or K, the category G in this 15 example, are marked with the positive indication at step 110. The remaining categories are marked with the negative indication (N) at step 112.

In the example shown in Figure 7c, TP=4 20 (categories A, G, H, K), FP=0 (as the category F is not counted because it was marked indifferent), TN=5 (categories C,D,E, I,J), and FN=1 (category B). Note that category B, which is marked negative (N) and true (T), is a "false negative" FN, as the classifier 25 12 incorrectly classified it as negative. In this example, $OC = 1/(4+1) = 0.2$ and $OA = 0/(0+4) = 0$. If the accuracy definition that uses the average of OC and OA is applied, accuracy = $1 - (0.2 + 0)/2 = 0.9$. If the harmonic mean definition is applied, accuracy 30 = 1 (because OA = 0).

Figure 8 illustrates the training of the classifier 12 in response to the accuracy 18

according to the present teachings. The accuracy 18
is provided to a modifier 20. The accuracy 18 may be
provided to the modifier 20 in the raw form of the
TP, FP, FN, and TN counts or may be in a combined
5 form.

The modifier 20 determines an alteration to be
applied to the classifier 12 in response to the
accuracy 18. The type of alteration applied to the
10 classifier 12 depends on the nature of the classifier
12. For example, if the classifier 12 is a rule-
based system then the alteration may be an alteration
to one or more of its rules. If, for example, the
classifier 12 is a neural network then the alteration
15 may be an alteration to neural network weights.

The training of the classifier 12 may be an
iterative process in which an alteration is applied
to the classifier 12 in response to the accuracy 18
20 for a test item and then in response to the accuracy
18 obtained for a next test item, etc.

The accuracy 18 may be used as a fitness measure
for designing the classifier 12 using genetic
25 programming techniques.

The present techniques for determining the
accuracy of a classifier may be used to analyze the
purchasing behavior of individuals. For example, the
30 authoritative classifier 14 may correspond to the
actual purchasing behavior of a customer and the
classifier 12 may correspond to the predicted
purchasing behavior according to a model or profiling

engine. The accuracy measure 18 may be used to improve the model and/or evaluate the model.

5 The present techniques for determining the accuracy of a classifier may be used to compare the performance of multiple categorization tools which may be provided, for example, by competing vendors.

10 The present techniques for determining the accuracy of a classifier may be used in an information management system such as a content, knowledge, or document management system. For example, the present techniques may be used to evaluate and/or improve the categorization of items 15 in a content, knowledge, or document management system.

20 The present techniques for determining the accuracy of a classifier may be used to evaluate and/or improve the placement of items in electronic portals - for example, web site links or descriptions of documents.

25 The present techniques may be applied to the categorization of products, goods, or services.

30 The foregoing detailed description of the present invention is provided for the purposes of illustration and is not intended to be exhaustive or to limit the invention to the precise embodiment disclosed. Accordingly, the scope of the present invention is defined by the appended claims.